Noname manuscript No.
(will be inserted by the editor)

A comparison of Hake's g and Cohen's *d* for analyzing gains on concept inventories

Jayson M. Nissen · Robert M. Talbot · Amreen Nasim Thompson · Ben Van Dusen

arXiv:1612.09180v1 [physics.ed-ph] 29 Dec 2016

Received: date / Accepted: date

Abstract Measuring student learning is a complicated but necessary task for understanding the effectiveness of instruction and issues of equity in college STEM courses. Our investigation focused on the implications on claims about student learning that result from choosing between one of two commonly used methods for analyzing shifts in concept inventories. The methods are: Hake's gain (**g**), which is the most common method used in physics education research and other discipline based education research fields, and Cohen's *d*, which is broadly used in education research and many other fields. Data for the analyses came from the Learning Assistant Supported Student Outcomes (LASSO) database and included test scores from 4,551 students on physics, chemistry, biology, and math concept inventories from 89 courses at 17 institutions from across the United States. We compared the two methods across all of the concept inventories. The results showed that the two methods led to different inferences about student learning and equity due to **g** being biased in favor of high pretest populations. Recommendations for the analysis and reporting of findings on student learning data are included.

Keywords Concept Inventories · Normalized Learning Gain · Effect Size · Student Learning · STEM Education

1 Introduction

The methods for measuring change or growth and interpretations of results have been hotly discussed in the research literature for over 50 years (Bereiter, 1963). Indeed, the idea of simply measuring a single state (let alone change) in an individual's understanding of a concept, conceptualized as a latent construct, is wrought with issues both philosophical and statistical (Cronbach and Furby, 1970). Despite these unresolved issues, measurement of growth has been a central tool for quantifying the effectiveness of interventions, treatments, and most importantly in our field, innovations in teaching and learning. Reporting gain scores and change

Jayson Nissen
E-mail: jnissen1@csuchico.edu
CSU Chico, Department of Science Education, 101 Holt Hall, Chico, CA 95929

metrics, often referenced against normative or control data, is commonplace and serves a strong basis for judging the efficacy of innovations. As measuring change and reporting gains and effects is common practice, it is incumbent on researchers to do so in the most accurate and informative manner possible.

In our present work, we harnessed the power of existing normative data and assessments, along with data collected at scale, to examine several statistical issues with characterizing change or gain in student understanding. The focus of our analyses in this investigation are on student scores on science Concept Inventories (CIs). CIs are research-based instruments that target common student ideas or prior conceptions. These instruments are most often constrained response (multiple choice) and include these common ideas as attractive distractors from the correct response. There exist a multitude of CIs in use across Biology, Chemistry, and Physics (our target disciplines) as well as in other fields (Engineering, Math). But just because so many CIs exist does not mean they are all *good* instruments. While CIs are common, the amount of validity evidence for each instrument varies widely, and normative data does not always exist for each CI. All of the CIs used in our work have at least some validity evidence and they are aligned with our proposed uses for the scores.

A principal tool in quantitative research is comparison, which leads to the frequent need for comparing across different instruments and contexts. These cross contextual comparisons are complicated by the fact that the instruments employed may have different scales and that the scores may greatly vary between populations. For example, some CIs are designed to measure learning across one semester while others are designed to measure learning across several years. These instruments could both be used in the same course but would, by design, provide very different results. To compare the changes on the two instruments it would be necessary to standardize the change in the scores, which is commonly done by dividing the change by a standardization coefficient. In much of the social science and education research this standardization coefficient is a measure of the variance of the scores.

The most common gain measurement used in Discipline Based Education Research (DBER) with CI data is the average normalized gain, $\mathbf{g}$, shown in equation 1 (Hake, 1998). In this equation the standardization coefficient is the maximum change that could occur in the course given the pretest mean. This standardization coefficient was adopted because it accounted for the smaller shift in means that could occur in courses with higher pretest means. Hake argued (p. 66) that $\mathbf{g}$ was a suitable measure because it was not correlated with pretest means, whereas posttest mean and absolute gain were correlated with pretest means and were not suitable measures. He also argued that this normalization allowed for "a consistent analysis over diverse student populations with widely varying initial knowledge states" because: the courses with lecture-based instruction *all* had low $\mathbf{g}$, courses with active-engagement instruction primarily had medium $\mathbf{g}$, and no courses had high $\mathbf{g}$. Hake then used this reasoning to define high $\mathbf{g}$ ($\mathbf{g} > 0.7$), medium $\mathbf{g}$ courses ($0.7 > \mathbf{g} > 0.3$), and low $\mathbf{g}$ courses ($\mathbf{g} < 0.3$).

$$\mathbf{g} = \frac{\bar{x}_{post} - \bar{x}_{pre}}{100\% - \bar{x}_{pre}} \tag{1}$$

Since Hake published his 1998 paper using $\mathbf{g}$, it has been very widely used with over 3,800 citations tracked in Google Scholar. Accordingly, there exists a

large amount of gain data expressed in terms of **g** in the research literature, all of which serves as normative data for other studies. While the use of **g** does not align with the practices of the broader social science fields, it would be naive to dismiss this metric as unimportant. However, as noted above and discussed in more detail below, some issues exist with **g**.

The field of education research has primarily adopted the effect size metric as the prefered method for measuring change. Perhaps the most commonly used effect size metric is Cohen's d (Cohen, 1977). In effect size metrics, a measure of the variance in the distribution of scores is used as the standardization coefficient rather than the maximum possible gain, which **g** uses. For our application, Cohen's d is given by equation 2, where the standardization coefficient s is the pooled standard deviation of the pre- and posttests (discussed further below).

$$d = \frac{\bar{x}_{post} - \bar{x}_{pre}}{s} \tag{2}$$

Extensive research into the utility and limitations of d has been performed, while the research investigating **g** is more limited. Coletta and Phillips (2005), in contrast to Hake's (1998) earlier finding, found that **g** was correlated with pretest means. Willoughby and Metz (2009) found that inferences based on **g** suggested gender inequities in college STEM courses even though several other measures indicated that there were not gender inequities in those courses. Furthermore, there are several different methods for calculating **g** that are commonly used and can lead to discrepant findings (Bao, 2006; Marx and Cummings, 2007). Issues have also been found for d, which is known to exaggerate the size of differences particularly when measuring changes in small samples (Grissom and Kim, 2012). d is based on the t-statistic and the assumptions of normality and homoscedasticity in the test scores that are used to generate it (Grissom and Kim, 2012). d's reliance on the assumptions of normality and homoscedasticity are often not met with CI data because of floor and ceiling effects as well as outliers. We expect that any problems that this creates for d are also applicable to **g**. However, we are not aware of any research on these assumptions pertaining to **g**.

2 Purpose

Both d and **g** are known to have limitations. Our purpose in this investigation was to empirically compare concept inventory gains using both **g** and d to investigate the extent to which the different standardization coefficients could lead to different inferences about student learning. In particular, we were concerned that bias in either measure could lead to incorrect inferences around issues of student equity. Our investigation was motivated by **g** being the de facto measure of student learning in physics education research and its adoption by other Discipline Based Education Research (DBER) fields despite there being very few investigations of the validity of **g** and known problems with the efficacy of **g**. We chose to compare **g** to d since d is gaining use in DBER and it is the comparable de facto measure in the much larger fields of sociology, psychology, and education research where its validity, utility, and limitations have been extensively studied.

3 Background on Measuring Change

In this section we provide a foundation for our motivations and work. We first discuss the development, adoption, and use of CIs in undergraduate science education research. We then discuss some statistical issues related to measurement of change before reviewing the uses of the average normalized gain in analyzing scores from CIs. Finally, we discuss the effect size metric and its use in the context of best practices for presenting data and findings.

3.1 Rise in the use of CIs to measure student knowledge

CIs provide "... data for evaluating and comparing the effectiveness of instruction at all levels." Hestenes and Wells (1992, p. 159). They typically consist of banks of multiple-choice items written to assess student understanding of canonical concepts in the sciences, mathematics, and engineering. CIs are generally developed through an iterative process of identifying core concepts with expert feedback and using student interviews to identify common preconceptions and provide wording for distractors. CIs exist for core concepts in most STEM fields, see Smith and Tanner (2010) for a thorough review and discussion. Though it is unclear exactly how many CI's currently exist, one of the most widely used CIs, the Force Concept Inventory (Hestenes et al, 1992), has been cited more than 2,900 times according to Google Scholar.

CIs are often used as the outcome measures for evaluative studies to find out if an instructional intervention has had an effect on learning relative to a control condition. To facilitate this, the CI items are administered pre and postinstruction, and gains observed in treatment groups are usually compared to gains observed in a control condition. CIs are designed to measure conceptual understanding at a "big picture" level. This means that if students conceive of science learning as a matter of primarily memorizing definitions and formulas (consistent with a more 'traditional' conception of teaching and learning), they are unlikely to do well on most CIs. Several studies have used CIs to compare the impact of research-based pedagogies to more traditional pedagogies (Hake, 1998; Freeman et al, 2014; Von Korff et al, 2016) and to investigate equity in STEM courses by comparing knowledge and learning of majority and underrepresented minority students (Madsen et al, 2013; Brewe et al, 2010; Rodriguez et al, 2012; Lorenzo et al, 2006).

Because scores are often compared for different CIs administered to different populations a change metric that is standardized and free from the original scale of the measurement is often used. This change metric is often $\mathbf{g}$ for DBER studies, but some studies have used d. One particular case that focused our current investigation on comparing $\mathbf{g}$ and d, was the use of $\mathbf{g}$ and test means to conclude, "In the most interactively taught courses, the pre-instruction gender gap was gone by the end of the semester", (Lorenzo et al, 2006, p. 1). A finding that was later called into question when an, "analysis of effect sizes showed gender still impacted FCI scores and that the effect was never completely eliminated", (Rodriguez et al, 2012, p. 6).

3.2 Some issues in measuring change

Discussions in the measurement literature on quantifying change can be sobering to say the least. A classical and often cited work in this area is that of Cronbach and Furby (1970), which raised issues of both the reliability and validity of gain scores. Based on Classical Test Theory, the authors argue that the prime issue of reliability has to do with the systematic relationship between error components of true scores derived from independent, but "linked" observations. Consider a common situation in CI use in which the same test form is given as both a pre- and posttest. One could argue that the observations (pre and post) are independent measurements since they are taken at different time points, but they are actually *linked* since the measurements are from the same group of students. Because those students had responded to the same instrument at the pretest administration, their posttest scores are likely to be correlated with their pretest scores due in large part to a shared error component between the two scores. One can correct for this (often overstated) correlation due to the shared error components, but the correction is not always straightforward. Bereiter (1963) calls this the "over correction under correction dilemma." Cronbach and Furby discuss this dilemma at length when decomposing the reliability of gain scores, and offer various methods to dissattenuate the correlation. However, the authors seem to see these correction methods as a workaround for the real issue of linked observations. In their summary discussion, Cronbach and Furby actually state that "investigators who ask questions regarding gain scores would ordinarily be better advised to frame their questions in other ways" (p. 80). Despite their persistent statistical issues, gain measurements are widely used in education research due to their great utility. Acknowledging these issues while leveraging the utility requires researchers to be diligent and transparent in their methods and presentation.

Another issue with gain scores has to do with the actual scale of the scores, which Bereiter refers to as the "physicalism-subjectivism dilemma." The issue here is related to the assumption of an interval scale on the construct of interest when using raw scores, or when using gain scores which are normalized on that same scale (as in using **g**). In other words, the gain metric (**g**) is scaled in terms of the measure itself (e.g., Newtonian Thinking as measured by a CI) and is assumed to be intervally scaled. A potential solution here is to change the scaling to something that "seems to conform to some underlying psychological units", (Bereiter, 1963, p. 5). In this case, the scaling factor (or "standardization coefficient") is not based in the scale of the measure (e.g., raw scores on a CI, or Newtonian Thinking) but rather in a "standard" unit such as variance of the score distribution. In this way the gain metric is transformed out of the scale of the measure (e.g., Newtonian Thinking) and into a construct independent, standardized scale (e.g., based on variance). The transformation of scale can make cross-scale comparisons possible, and also may highlight potential inequities brought on by remaining in the scale of the measure itself. This latter approach is how the dilemma is addressed when using the effect size metric, discussed further below. For a more detailed discussion of these issues related to the measurement of change in Classical Test Theory see (Author).

3.3 The Average Normalized Gain

The average normalized gain (**g**) was developed by Hake as a way to normalize average gain scores in terms of how much gain could have been realized (Hake, 1998). Hake interpreted **g** from pre-post testing "as a rough measure of the effectiveness of a course in promoting conceptual understanding" (p. 5). His work was seminal in physics education research and has lead to the broad use of **g** throughout DBER. This broad use has lead to at least three different methods for calculating **g** to currently be in common use. The original method proposed by Hake calculates **g** from the group means and is shown in equation 1. A second method that is more commonly used (Von Korff et al, 2016) is to calculate the normalized gain for each individual student ($\mathbf{g}_I$) to characterize that student's growth, and to then average the normalized gains for all of individuals to calculate **g** for the group. Bao (2006) provides an in depth discussion of the affordances of these two methods, but Bao and Hake both state that in almost all cases the two values are within 5% of one another.

Marx and Cummings (2007) proposed a third method (normalized change, **c**) in response to a number of shortcomings with respect to **g** and its use that they noted. The issues they note with **g** included: a bias towards low pretest scores, a non-symmetric range of scores ($-\infty$ to 1), and a value of -∞ for any posttest score when the student achieves a perfect pretest score. These limitations inhibit the ability to calculate the average normalized gain for a class by averaging the individual student normalized gains. Instead, they offer a set of rules for determining **c** based on whether or not the student gains from pre- to posttest, worsens, or remains at the same score. Their metric results in values of **c** ranging from -1 to +1. However, **c** still suffers from a bias towards low pretest scores as noted by Bao, and is still normalized on the same scale as the measure itself, an issue raised by Bereiter (1963) and discussed above.

One particular concern with gain metrics, and with **g** specifically, has to do with the possibility that these metrics can be biased for or against different groups of students. Willoughby and Metz (2009, p. 1) found that "males had higher learning gains than female students only when the normalized gain measure was utilized. No differences were found with any other measures, including other gain calculations, overall course grades, or individual exams." One might expect this to be the case when pretest score is part of the standardization coefficient, since pretest is likely correlated with previous education, and therefore opportunity and even socioeconomic status. Indeed, Coletta and Phillips (2005, p. 1) "found a significant, positive correlation between class average normalized FCI gains and class average preinstruction scores." This finding is aligned with Marx and Cummings (2007) conclusion that **g** is biased by pretest scores, however, they found it was biased the opposite direction.

3.4 The Effect Size Metric

One of the mostly widely used standardized effect size metric is Cohen's *d*. Cohen's *d* normalizes (i.e. "scales") the difference in scores in terms of the standard deviation of the observed measurements. In essence, it is the difference between Z (standard) scores. This results in a "pure" number that is free from the original

scale of measurement (Cohen, 1977). As a result, d meets the need for "... a measure of effect size that places different dependent variable measures on the same scale so that results from studies that use different measures can be compared or combined", Grissom and Kim (2012).

As a consequence of using the standard deviation, d assumes that the populations being compared are normally distributed and have equal variances. Accordingly, the standard deviation used to calculate d is that of either sample from the population since they are assumed to be equal. However, in practical applications the pooled standard deviation of the two samples is used since they often vary slightly. In our work we use equation 3 to calculate the pooled standard deviation (s) for pre- and posttests.

$$s = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}} \tag{3}$$

Using either the equal pre- and posttest standard deviations or the pooled standard deviations assumes that the samples (pre- and posttest) are independent, and therefore does not take into account or correct for the correlation between measurements made at pre and post (the "dilemma" discussed above from Bereiter). d can also be directly calculated from the t-test statistic and this can serve to further elucidate the dependent-independent issue. Dunlap et al (1996) present an example of calculating a t-statistic between two means when assuming the samples are independent, and again when assuming dependence for the same sample means (correlated at 0.6857). The dependent t-statistic (4.513) is much larger than the independent (2.530), as the "correlation between measures reduces the standard error of the difference between the means, making the difference across conditions more identifiable" (p. 171). In other words, the degree of the difference between the means (as expressed by the magnitude of the t-statistic) is overstated when the scores are correlated. Therefore, if d is calculated assuming independent samples (when the samples are actually dependent) the degree of the effect *will be overstated if it does not take into account the correlation between the two measurements* (Cohen, 1988). In practice, many researchers use equation 3 and do not account for correlations between pre- and posttest. This is an acceptable practice so long as all researchers are aware of the issue and any effect sizes being compared are calculated in the same way (Morris and DeShon, 2002).

4 Research Questions

Given our purpose of comparing **g** and d, our specific research questions were: 1) What were the relationships between each of these gain metrics? 2) To what extend did the different measures of change lead to different inferences about certain groups of students?, and (3) How robust were against **g** and d outlier data? Based on previous research we expected some differences in the degree to which d and **g** indicated that a phenomenon (i.e. gain) was present. We expected the gain characterization by each metric to vary by student population due to differences in pretest scores across populations. This latter point motivated our research because it could bias investigations of equity in college STEM learning.

5 Methods

In order to answer these research questions we used a large data set of student responses to nine different research-based CIs. This data set was large enough to provide useful and reliable comparisons of effect size measures and to represent CI data in general. We processed the data to remove spurious and unreliable data points and used Multiple Imputations (MI) to replace missing data. In order to simplify our analysis, we first investigated the similarity of the measures resulting from the three ways to calculate Hake's normalized gain, course averages ($\mathbf{g}$), averaged individual gains($\mathbf{g_I}$), and normalized change ($\mathbf{c}$), to determine if they were similar enough that we could conduct our further analysis using only one of those approaches. We then made several comparisons of the effect size measures for $\mathbf{g}$ and d to inform our research questions. We compared the effect size measures to one another and investigated the relationships of the effect size measures with pretest and posttest means and standard deviations to identify any potential bias in the effect size measures. To test any biases we found and to inform the effects of those biases, we compared the effect size measures for subpopulations within each course that have historically different pretest and posttest means. We also compared the effects of outlier data on each of the effect size measures to inform the robustness of each measure.

5.1 Data collection and processing

Data for the study was collected as part of a project to assess the impact of Learning Assistants (LAs) on student learning (Otero, 2006). LAs are talented undergraduates who are hired by university and two-year college faculty to help transform courses. To facilitate the collection of student learning data from across many institutions using LAs, the LA Supported Student Outcomes (LASSO) online assessment platform was created. LASSO is a free platform that is hosted on the LA Alliance website (https://www.learningassistantalliance.org/) and allows faculty (LA-using or not) to easily administer research-based concept inventories as pre- and posttests to their students online. To use LASSO, faculty provide course-level information, select their assessment(s), and upload a list of student names and emails. When faculty launch an assessment their students receive emails with unique links to complete their tests online at the beginning and end of instruction. Faculty can track students' participation and send reminder emails. As part of completing the instrument, students answer a set of demographic questions. Faculty can download their students' responses as well as a summary report that shows the distribution of their students' pre- and posttest scores, normalized learning gains ($\mathbf{g}$), and effect sizes (Cohen's d).

We processed the data from LASSO to ensure that spurious data points were removed and courses had sufficient data for reliable measurements. We included courses that had partial data for at least 10 participants in order to meet the need for a reliable measure of means without excluding small courses from our analyses. In order to remove spurious and unreliable data, student and course level data was removed if:

Table 1 Data after each filter was applied.

	None	Time	Completion	Gain	$\geq 60\%$ missing
Courses	119	116	116	116	89
Students	6,041	5,677	5,667	5,667	4,551
Pretests	5,339	4,922	4,899	4,899	3,842
Posttests	4,204	3,693	3,685	3,642	3,335
Matched	3,502	2,973	2,917	2,874	2,626

- A student had taken less than 5 minutes to complete that test. We reasoned that this was a minimum amount of time required to read and respond to the test questions.
- A student answered less than 80% of the questions on that test. We reasoned that these exams did not reflect student's actual knowledge.
- A student's absolute gain (posttest mean minus pretest mean) was 2 standard deviations below the mean absolute gain for that test. In these cases, we removed the posttest scores because we reasoned that it was improbable for students to complete a course and unlearn the material to that extent.
- A course had greater than 60% missing data. We were concerned that low response rates indicated abnormal barriers to participating in the data collection that could have influenced the representativeness of the data from those courses

Prior to our analysis we imputed test scores for students missing either the pretest or posttest using the Amelia II package in R in order to maximize the power and reliability of our analyses. We used multiple imputations because it reduces the rate of type I error by using all of the available information to better account for missing data (Rubin, 1996). This leads to much better analytics than traditional methods such as listwise deletion (Allison, 2002) that, while they "...have provided relatively simple solutions, they likely have also contributed to biased statistical estimates and misleading or false findings of statistical significance", (Buhi et al, 2008, p.400).

Our multiple imputation model included variables for pretest and posttest scores, duration to complete each test, dummy variables for each CI, participant's year in school, whether this was the participant's first time taking the course, and participant's belonging to an underrepresented minority group for race and for gender. The data collection system required or generated responses to all of these questions except for test duration and test score. Some data was missing for 42% of the students. Based on this 42% rate we conducted 42 imputations because this is a conservative number that will likely provide better results than a smaller number of imputations (Bodner, 2008). We analyzed all 42 imputed data sets and combined the results using Rubin's Rules (Schafer, 1999). All assumptions were satisfactorily met for all analyses. For the stability analysis, discussed below, we analyzed a single aggregated dataset that was formed by averaging the 42 imputed datasets.

5.2 Investigating the effect size measures

In order to identify and investigate differences between the effect size measures we used correlations and multiple linear regressions (MLR) to investigate the relationships between the effect size measures and the test means and standard deviations. Correlations informed the variables we included in the MLR models.

To measure the size of the impact that any bias in one of the the effect sizes would have on concept inventory data, we used the male and female effect size measures in the aggregated data set. We separated these two groups because male students tend to have higher pretest and posttest means on science concept inventories than female students (Madsen et al, 2013; Cunningham et al, 2015). Thus, gender provided a straightforward method of forming populations with systematically different test means and standard deviations. Gender also allowed us to frame our analysis in terms of equity of effects. We adopted a definition of equity being the case where a course does not increase pre-existing differences. This definition means that for a course to be equitable the effect on the low pretest group is equal to or larger than the effect on the high pretest group.

We applied this definition to our purpose by calculating the effect sizes for males and females separately. For each effect size measure we then calculated the difference between males and females effect sizes, for example $\Delta_{\mathbf{g}} = \mathbf{g}_{Male} - \mathbf{g}_{Female}$. If males in a course had a larger effect size than females in the course then that course was inequitable, $\Delta_{\mathbf{g}} > 0$. This created four categories into which any two effect size measures would locate each course. Two categories for agreement where both effect sizes said it was either equitable or inequitable and two categories for disagreement where one said equity and the other said inequity. If one of the effect size measures was biased and indicated larger effects on higher pretest mean populations then we expected that one type of disagreement would occur more frequently than the other type. To easily identify differences in the number of courses in the disagreement categories and the size of those disagreements, we plotted the data on a scatter plot. We tested the statistical difference in the distributions using a chi square test of independence with categories for each effect size measure and whether they indicated equity or inequity.

In order to test the stability of the both Cohen's d and Hake's $\mathbf{g}$ when spurious data are analyzed, we added outliers to the data set for each course. We repeated this process four times so that five data sets were generated: the original, plus one outlier, plus two outliers, etc. We then calculated d and $\mathbf{g}$ for each of these five data sets. The outliers were pretest scores of a mean for the course and posttest scores of 100%. We chose these outliers because they were realistic possibilities that represented a scenario where a student did not represent the effects of the course because they had exceptional ability or resources that were not available to other students. Because d and $\mathbf{g}$ are based in different scales we used the percentage change from the initial value to compare the effects for the addition of each outlier. We then used a boxplot of the percentage change for the effect sizes to compare their stability and repeated measures ANOVA to check the null hypothesis that the changes in $\mathbf{g}$ and d were not different from one another.

Table 2 Correlations between the three forms of normalized gain. *** indicates $p<0.001$

	c	**g**	$\mathbf{g_I}$
c	1.00***		
g	0.99***	1.00***	
$\mathbf{g_I}$	0.93***	0.93***	1.00***

5.3 Simplifying the analysis

The multiple methods for calculating normalized gain complicated our purpose of comparing normalized gain and Cohen's d. Therefore, we compared normalized gain calculated using the three methods described in the introduction for each course. Correlations between all three measures were all large and statistically reliable, as shown in Table 1. These results indicated that all three measures were very similar and we only used the normalized gain calculated using course averages, **g**, in our subsequent analyses.

6 Findings

6.1 Relationship between **g** and d

Investigating the relationship between **g** and d indicated that their was a large positive relationship between the two measures and it was statistically reliable, r=0.73, $p<0.001$. This indicated that d and **g** shared approximately half of their variance in common (adjusted r^2=0.53). Because these two measures serve the same purpose, the 47% that they *do not* have in common was a large amount. Further investigating the correlations between the effect sizes and their related descriptors, which are shown in Table 2, revealed large differences between d and **g**. The correlations between d and both pretest mean and pretest standard deviation were small to very small and were not statistically reliable. In contrast, **g** was moderately to strongly correlated with both pretest mean and pretest standard deviation. These correlations between **g** and pretest statistics indicated that approximately one fifth of the variance in normalized gains was accounted for by the score distributions that students had prior to instruction. In contrast, d was only weakly associated with both pretest mean and pretest standard deviation. This was strong evidence that **g** was positively biased in favor of populations with higher pretest scores.

To inform the size of this bias, we ran several models using MLR with **g** as the dependent variable and independent variables for d, pretest mean, and pretest standard deviation. We used **g** as the dependent variable because this was consistent with the correlations between **g** and pretest mean indicating that **g** was positively biased by pretest means; whereas correlations indicated that d was not biased. Our focus in these MLRs was on the additional variance that was explained by each variable in the model. All of the models were statistically significantly reliable ($p<0.05$). The first model only included d and shows that d and **g** shared 53% of the same variance as indicated by the adjusted r^2 value. The addition of either pretest mean or pretest standard deviation to the model markedly increased

Table 3 Correlations between the effect size measures and test statistics. * $p<0.05$. ** $p<0.01$. *** $p<0.001$.

	d	**g**	Pre. Mean	Post. Mean	Gain	Pre. S.D.	Post. S.D.
d	1.00***						
g	0.75***	1.00***					
Pre. Mean	-0.11	0.43***	1.00***				
Post. Mean	0.43***	0.87***	0.81***	1.00***			
Gain	0.87***	0.93***	0.11	0.67***	1.00***		
Pre. S.D.	-0.18	0.44***	0.67***	0.65***	0.24*	1.00***	
Post. S.D.	-0.24	0.10	-0.01	0.08	0.15	0.42***	1.00***

Table 4 MLR exploring relationships between **g**, *d*, and population statistics.

Model	1		2		3		4	
r^2(%)	55.6		82.4		89.0		92.1	
$adj.r^2$(%)	55.0		82.0		88.7		91.8	
p	<0.001		<0.001		<0.001		0.02	
D.V.	Beta	*p*	Beta	*p*	Beta	*p*	Beta	*p*
Intercept	0.05	0.05	-0.21	<0.001	-0.21	<0.001	-0.26	<0.001
d	0.22	<0.001	0.24	<0.001	0.25	<0.001	0.25	<0.001
Pre. Mean			0.01	<0.001			0.01	<0.001
Pre. S.D.					0.01	<0.001	0.01	<0.001

the explained variance to either 82% or 89%, as is shown in Table 3. Including all three variables explained 92% of the variance in **g**. We interpreted this as indicating that the disagreements between *d* and **g** were largely explained by the pretest means and standard deviations. Because the pretest means and standard deviations explaining such a large proportion of the unexplained variance from Model 1 and the correlations indicating that pretest mean and pretest standard deviation were much more strongly related to **g** than to *d*, these results indicated that **g** was biased in favor of groups with higher pretest means.

6.2 Testing the bias in **g** on historically different populations

In order to visualize the equity of effects in the courses we plotted the difference in *d* on the x axis and the difference in **g** on the y axis. We colored the course markers to indicate whether male or female students' pretest means were higher. The vast majority of markers (41 of 43 courses) indicated that male students started with higher pretest means and that the data was consistent with our focus on equity being indicated by a larger effect on female students. In total **g** indicated a larger effect on males in 33 out of 43 courses whereas *d* indicated a larger effect on males in 22 out of 43 courses. Figure 1 illustrates this difference clearly in the difference between Quadrants II and IV. A chi squared test of independences indicated that these differences were statistically reliable, $\chi^2(1) = 6.10$, p=0.013. This difference confirmed that **g** was biased in favor of populations with higher pretest means. This bias implies that **g** is not a sufficiently standardized change metric to allow

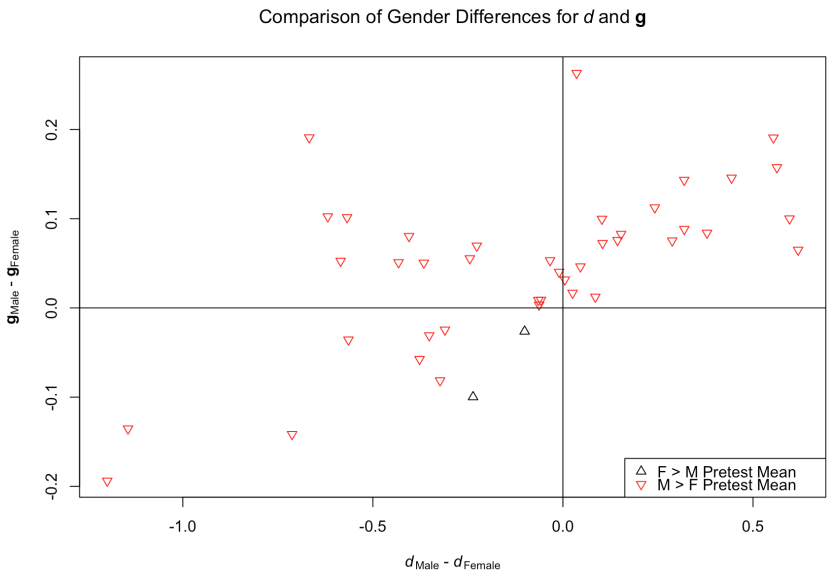

Fig. 1 Comparison of Gender Differences for d and **g**

comparisons across populations or instruments that have different pretest means and is not a suitable measure of effects.

6.3 Stability

We explored the stability of d and **g** through the addition of outliers to course data sets to identify how much each measure was affected by outlier data in the form of perfect posttest scores. Results indicated that outliers had a smaller effect on d than on **g**. In every course the percentage change for each outlier added was larger for **g** than for d. Figure 2 illustrates this difference by showing that the median percentage change across all of the courses for each additional outlier added was larger for **g** than for d. The first outlier added had a median effect on **g** of 8.6% and a median effect on d of 6.3%. A repeated measures ANOVA indicated that there was a statistically reliable difference in the change between d and **g** ($F(1,708)=13.3$, $p<0.001$) and while this difference increased as more outliers were added this increase was not statistically reliable ($F(1,708)= 2.98$, $p=0.08$). The greater change in **g** indicated that d was more robust against the effects of spurious data that can occur in CI datasets.

7 Discussion

Before comparing the statistical merits of using **g** and d to measure student learning, we had to determine what differences there were between varying methods of calculating **g**. Our analysis showed that the three methods for calculating normalized gain scores were all highly correlated ($r \geq 0.93$). The high level of correlation

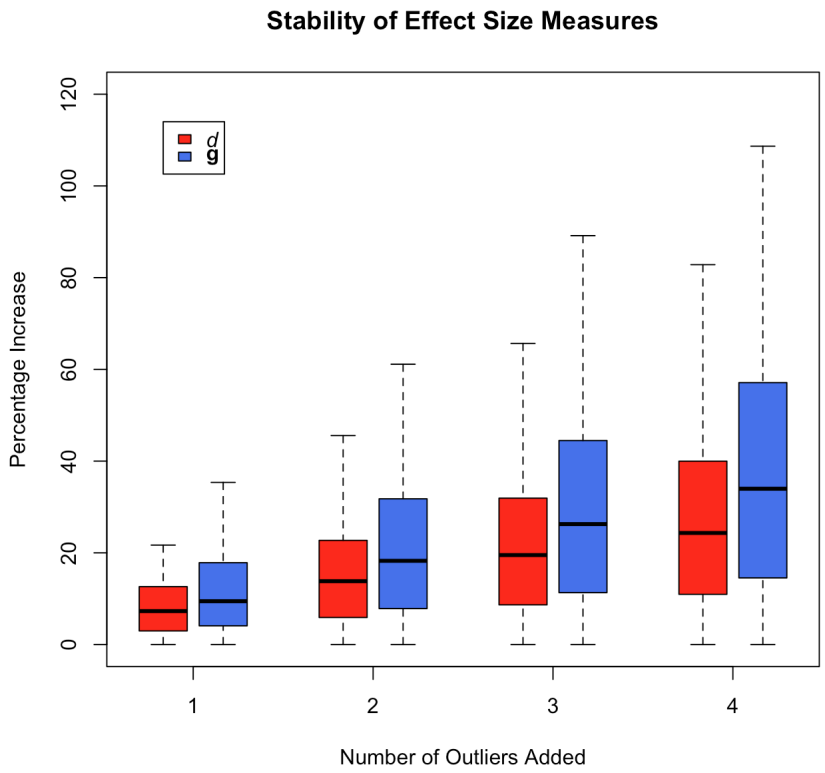

Fig. 2 Boxplot of the percent change in each course's effect size measure with the addition of outliers that had a pretest score of the course average and a posttest score of 100%. The outliers are not plotted.

between the normalized gain values indicates that it makes little difference which method is used. This is encouraging given that many publications report **g** scores without discussing which method of calculation was used (Von Korff et al, 2016). Because our analyses was focused on course level effects we used the course's average prescore to calculate **g** throughout our analyses.

In examining the stability of the two measures, d was shown to be more robust to outlier data than **g**. The addition of outlying student test data into a course was shown to have an approximately 50% larger impact on **g** than d. In addition to making d more representative of the average student than **g**, the robustness of d makes it a more reliable measure of gains on concept inventories which commonly include outliers and other issues with normality and heteroscedasticity. For both **g** and d the addition of a single outlier did not change the measure to the extent that different conclusions about the course would be drawn. However, in a few courses there were shifts in **g** and d that could lead to different inferences about the effectiveness of the courses. While these large shifts occurred more frequently for **g** than for d, users of these metrics should check the distributions of the populations they are measuring for outliers, normativity, and heteroscedasticity, particularly when the sample sizes are small.

Our comparisons of **g** and d revealed a number of meaningful differences that indicated that **g** was biased in favor of high pretest populations. The correlation between **g** and d was strong (r=0.73, $p<0.001$) but was markedly smaller than the correlations between the three different methods of calculating **g** ($r \geq 0.93$). This correlation of 0.73 means that **g** and d shared only 53% of their variance. MLRs indicated that most of the difference in **g** and d was explained by pretest mean and standard deviation; d, pretest mean, and pretest standard deviation accounted for 93% of the variance in **g**. Given that **g** was correlated with these pretest statistics much more strongly than d, our conclusion is that **g** is biased in favor of populations with high pretest means.

This bias of **g** in favor of populations with high pretest means is problematic. In fact, 21% of a class's **g** value was explained by the class' pretest mean. **g**'s dependence on pretest means privileges populations of students who come into a class with more disciplinary knowledge or who simply perform better on multiple choice exams. This bias disproportionately affects students from traditionally underrepresented backgrounds. When comparing the learning of males and females in our dataset **g** identified males as learning more in 33 of 43 courses (77%) while d only identified males as learning more in 23 of 43 courses (53%), nearly cutting the rate by 1/3 (Figure 1). This difference in measurement indicated that **g** should not be used for investigations of equity as it dramatically overestimated student inequities.

8 Recommendations for effect sizes on CIs

Our analyses show that using Cohen's d offers advantages over Hake's normalized learning gain (**g**) and is our generally recommended method for calculating student learning using pretest and posttest data. We also recommend that prior to calculating d researchers use multiple imputations to deal with missing data rather than listwise deletion. Researchers should include descriptive statistics to allow subsequent researchers to interpret and use the effect sizes in larger meta-analyses. These descriptives should include means, standard deviations, and sample sizes for each measure used, as well as correlations between the measures. We include correlations on this list because of the dependent nature of CI pre-post testing is not taken into account by the change metrics that we have presented in this paper. As discussed above in the background section, this correlation (i.e. "linking") results in a shared error component that exaggerates the size of the difference. While it is not a common practice in the field of education research, there are statistical methods that can account for the dependence of pre-post tests when the correlational values are known.

For *instructors* who wish to measure the effectiveness of their instruction but do not want to perform MI's we recommend not using listwise deletion and instead calculating d using all available data. This requires foregoing running a t-test on the course data to calculate a p-value due to the dependent nature of the data. If statistical hypothesis testing *is necessary* then MI should be used.

9 Conclusion and Inferences

When comparing the impact of interventions across settings and outcomes in terms of gain scores, some form of normalization must be employed. In this investigation we compared student learning across institutional settings, disciplines, and concept inventories. Normalized learning gain (**g**) and Cohen's d both employ standardization coefficients to try to account for the inherent differences in the data. If these standardization coefficients are biased, they skew the results. **g** was developed to account for classes with higher pretest means potentially having lower gains due to ceiling effects. By focusing on ceiling effects, **g** implicitly assumes that any population with a higher pretest score will have more difficulty in making gains than lower pretest populations. This assumption contradicts one of the most well established relationships in education research; that prior achievement is a strong predictor of future achievement. Thus, **g**'s adjustment for potential ceiling effects appears to overcorrect for the problem and results in **g** being biased in favor of populations with higher pretest means.

This systematic bias of **g** has implications for students and instructors. Students from traditionally underrepresented backgrounds are disproportionately likely to have lower pretest scores, thereby putting them at a disadvantage when making instructional or curricular decisions about intervention efficacy based on **g**. Instructors who teach courses (e.g. non-major courses) or are at institutions (e.g. 2-year colleges) that serve students who tend to have lower pretest means are disadvantaged when using **g** to measure their students' learning. This is particularly important for faculty at teaching intensive institutions where evidence of student learning can be an important criteria for tenure and promotion.

The bias of **g** is also an issue for researchers who wish to measure the impact of interventions on student learning. The efficacy of interventions ranging from curricular designs to classroom technologies have been evaluated and scaled-up based on measures of student learning. For all of these investigations, it is important to have a measure of student learning that is not significantly dependent on the knowledge that students bring to a class. By using the pooled standard deviation, rather than the maximum possible gain as defined by the pretest, as a standardization coefficient, d avoids the bias toward higher pretest means while simultaneously providing a score that is more robust against outliers and accounts for instrument specific difficulty of improving a raw score. For these reasons we recommend the use of d rather than **g** for measuring student learning. In addition to being the more reliable statistical method for calculating student learning, the use of d by the DBER community would align with those of the larger education research community, thereby facilitating greater cross-disciplinary conversation and collaboration.

References

Allison PD (2002) Missing Data. Sage, Thousand Oaks, CA

Bao L (2006) Theoretical comparison of average normalized gain. American Journal of Physics 74(10):917–922

Bereiter C (1963) Some persisting dilemmas in the measurement of change. In: Harris CW (ed) Problems in the measurement of change., University of Wisconsin Press

Bodner TE (2008) What improves with increased missing data imputations? Structural Equation Modeling: A Multidisciplinary Journal 15(February):651–675

Brewe E, Sawtelle V, Kramer LH, O'Brien GE, Rodriguez I, Pamela P (2010) Toward equity through participation in modeling instruction in introductory university physics. Physical Review Special Topics Physics Education Research 6(1)

Buhi ER, Goodson P, Neilands TB (2008) Out of sight, not out of mind: Strategies for handling missing data. American Journal of Health Behavior 32(1):83–92

Cohen J (1977) Statistical power analysis for the behavioral sciences. New York: Academic Press

Cohen J (1988) Statistical power analysis for the social sciences. Hillsdale, NJ: Erlbaum

Coletta VP, Phillips JA (2005) Interpreting FCI scores: Normalized gain, preinstruction scores, and scientific reasoning ability. American Journal of Physics 73(12):1172–1182

Cronbach LJ, Furby L (1970) How should we measure "change"- or should we? Psychological Bulletin 74(1):68–80

Cunningham BC, Hoyer KM, Sparks D (2015) Gender differences in science, technology, engineering, and mathematics (STEM) interest, credits earned, and NAEP performance in the 12th grade. National Center for Education Statistics

Dunlap WP, Cortina JM, Vaslow JB, Burke MJ (1996) Meta-analysis of experiments with matched groups or repeated measures designs. Psychological Methods 1(2):170

Freeman S, Eddy SL, McDonough M, Smith MK, Okoroafor N, Jordt H, Wenderoth MP (2014) Active learning increases student performance in science, engineering, and mathematics. Proceedings of the National Academy of Sciences U S A

Grissom RJ, Kim JJ (2012) Effect Sizes for Research: Univariate and Multivariate Applications, Second Edition. Routledge

Hake R (1998) Interactive-engagement vs. traditional methods: a six-thousand-student survey of mechanics test data for introductory physics courses. American Journal of Physics 66(1):64–74

Hestenes D, Wells M (1992) A mechanics baseline test. Phys Teach 30(3):159–166

Hestenes D, Wells M, Swackhamer G (1992) Force concept inventory. Physics Teacher 30(3):141–158

Lorenzo M, Crouch CH, Mazur E (2006) Reducing the gender gap in the physics classroom. American Journal of Physics 74(2):118 ST–Reducing the gender gap in the physics c

Madsen A, McKagan SB, Sayre EC (2013) Gender gap on concept inventories in physics: What is consistent, what is inconsistent, and what factors influence the gap? Physical Review Special Topics Physics Education Research 9(2):020,121

Marx JD, Cummings K (2007) Normalized change. American Journal of Physics 75:87–91

Morris SB, DeShon RP (2002) Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. Psychological Methods

7(1):105–125
Otero V (2006) The colorado learning assistant model: A multidisciplinary approach to teacher recruitment and preparation. Bulletin of the American Physical Society
Rodriguez I, Brewe E, Sawtelle V, Kramer LH (2012) Impact of equity models and statistical measures on interpretations of educational reform. Physical Review Special Topics Physics Education Research 8(2)
Rubin DB (1996) Multiple imputation after 18+ years. Journal of the American Statistical Association 91(434):473–489
Schafer JL (1999) Multiple imputation: A primer. Statistical Methods in Medical Research 8(1):3–15
Smith JI, Tanner K (2010) The problem of revealing how students think: Concept inventories and beyond. CBE Life Science Education 9(1):1–5
Von Korff J, Archibeque B, Gomez KA, Heckendorf T, McKagan SB, Sayre EC, Schenk EW, Shepherd C, Sorell L (2016) Secondary analysis of teaching methods in introductory physics: A 50k-student study. American Journal of Physics 84(12):969–974
Willoughby SD, Metz A (2009) Exploring gender differences with different gain calculations in astronomy and biology. American Journal of Physics 77(7):651–657